ATTACHMENT 3

WHITEPAPER

OVERTURE

CARRIER ETHERNET FOR CLOUD COMPUTING

The virtualized services model of cloud computing has significant implications on today's network, and how it is evolving to the network of tomorrow. The network impact of cloud computing is in fact so significant that any discussion on cloud computing should also include cloud networking.

Without the network, there could be no cloud. Since Carrier Ethernet is the optimal networking solution for connection oriented services and the optimal Layer-2 for IP services, is also the optimal solution for cloud computing.

The best way to see the role of Carrier Ethernet in cloud computing is to look at the similarities and differences in network infrastructure as applications and services evolve from a traditional enterprise-centric architecture to a datacenter enabled architecture and finally to a cloud computing architecture.

> Cloud Computing cannot exist without Cloud Networking. Carrier Ethernet is the optimal solution for cloud access and network connectivity within the cloud.
>
> — Scott Knox, Director Solutions Development

TRADITIONAL ENTERPRISE ARCHITECTURE

The traditional enterprise network architecture has most, if not all, of the applications, computing resources and databases located at the enterprise location (Figure 1). The CAPEX and OPEX of the servers and applications are dedicated to and managed by he enterprise.

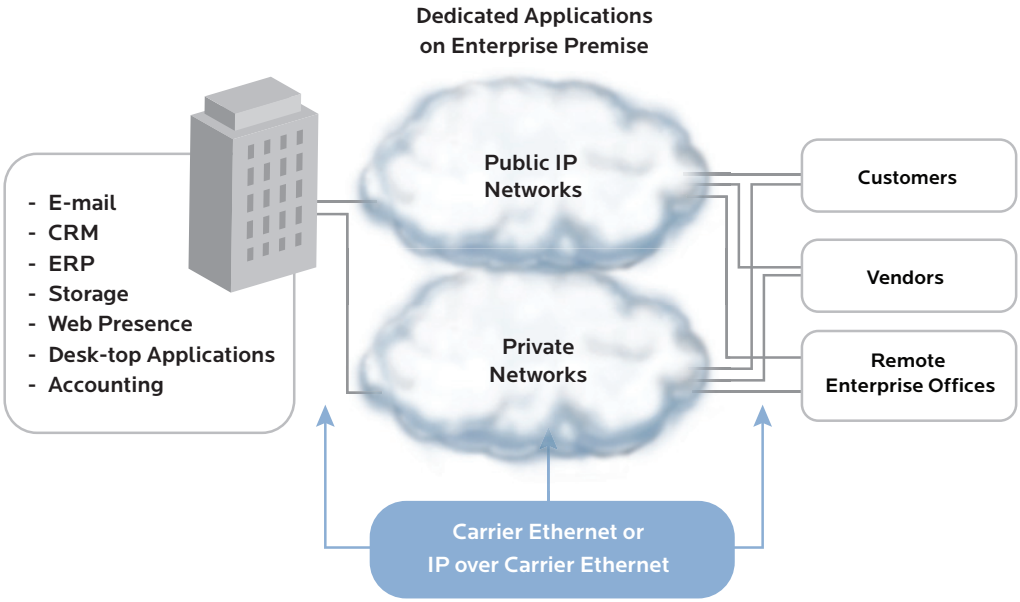

Figure 1. Carrier Ethernet in Traditional Enterprise Architecture

Network Connectivity to remote enterprise offices, customers, vendors and partners is via a combination of public and private networking. There tends to be many connections via private networking to dedicated users and applications, and a few large connections to applications that are driving scalable bandwidth. The connection to the public network

is typically a broadband connection for user access to public Internet content, and is also sometimes used with Layer-3 VPN solutions to provide secure tunnel connections for intra-net applications.

Attributes of traditional enterprise networks include:

- On-site users have Local Area Network (LAN) connectivity to applications.
- Multi-location enterprises use private networking to connect remote offices to applications, or have separate dedicated applications at each location
- Enabling applications for mobile users can be difficult
- Access bandwidth and external connection needs can scale quickly for applications such as web presence, Customer-to-Business (CTB) or Business-to-Business (BTB) e-commerce.

Carrier Ethernet enables CAPEX and OPEX efficiencies to all aspects of the end-to-end service for traditional enterprise networks including:

- End-end private networking with Layer-2 Ethernet
- Ethernet Access to Public-IP and Private-IP networks
- Layer-1 and/or Layer-2 infrastructure for Public-IP and Private-IP networks

DATACENTER ENABLED ARCHITECTURE

Datacenters offer the economy of large-scale, purpose-built space, power, HVAC (Heating, Ventilation and Air Conditioning) and management of computers, servers and storage. Fully equipped server space costs $400-$600 per square foot; outsourcing the construction and management to datacenter management companies is more cost effective and removes the direct CAPEX and OPEX burden from enterprises. An added bonus is a high degree of scalable, resilient network connectivity from multiple network providers.

The typical enterprise today has moved at least some portion of their applications into a datacenter environment (Figure 2). The first applications to move to the datacenter are those that require a large-scale server environment, and those that require a large amount of bandwidth and connectivity outside the enterprise such as web hosting. Other applications such as storage and backup must be moved off-site to meet redundancy and Sarbanes-Oxley requirements. Additional applications move depending on the scale and scope of the enterprise and their IT management model.

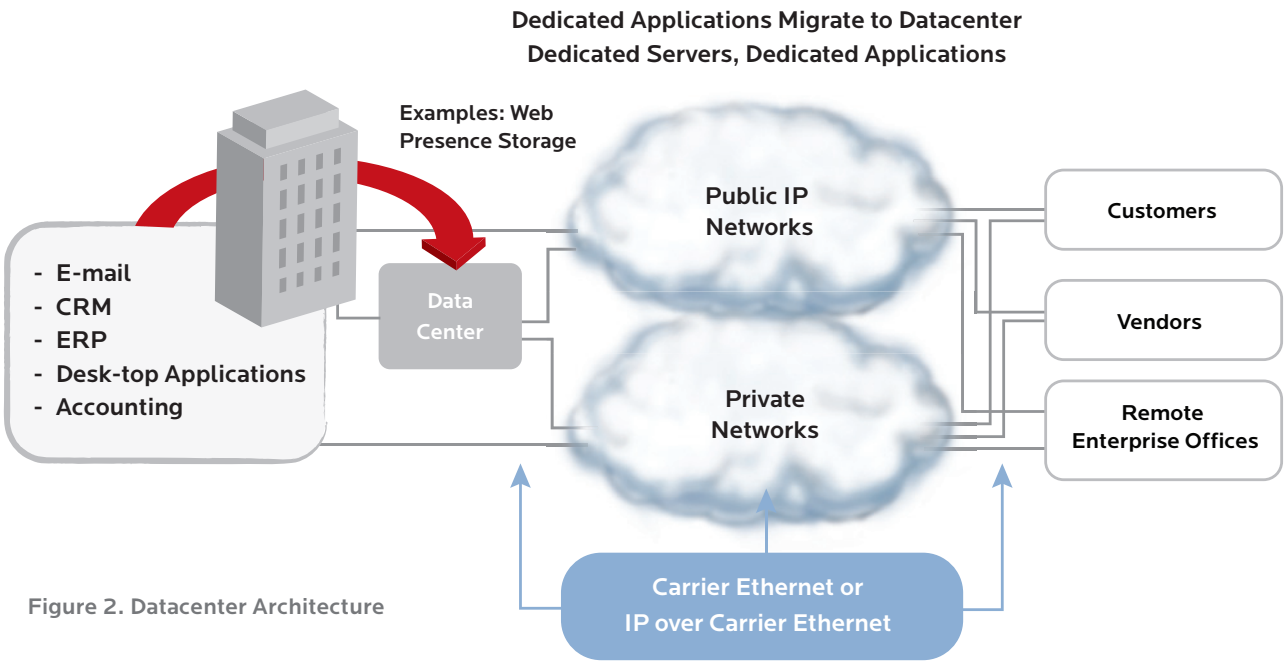

Figure 2. Datacenter Architecture

The business relationship between the enterprise and the datacenter service provider can range from basic co-location to fully managed and hosted services. In the co-location model, the enterprise owns and manages the applications on the servers; with the datacenter providing the space, power, and connectivity.

In the fully managed model, the datacenter owns and manages the servers as well, with servers typically dedicated for each enterprise. There are also solutions where the datacenter owner provides physical management of the servers, and a 3rd party ASP (Applications Service Provider) provides the applications and high-touch application management.

Datacenters change the networking profile in several ways. For applications which have significant connectivity outside the enterprise such as web presence, this reduces the need for access bandwidth to the enterprise since users connect via the Internet directly to the web servers at the datacenter. On the other hand, access bandwidth for applications internal to the enterprise increases when these move outside the enterprise. The scaling bandwidth needs for all applications means a net significant increase in total networking capacity.

Most enterprises make significant changes and upgrades to their IT architecture when applications and servers are transitioned to a datacenter. Instead of simply moving 'old iron' to the datacenter, most enterprises upgrade to new server platforms and migrate from legacy network technologies such as ATM, Frame Relay, or Fiber Distributed Data Interface (FDDI) to Carrier Ethernet.

The net impact of datacenters on the network architecture is a dramatic increase in the need for Carrier Ethernet solutions. The quantity of bandwidth, the increase in performance and reliability requirements and the focus on CAPEX and OPEX efficiency make the move the Carrier Ethernet the optimal choice.

CLOUD COMPUTING

In cloud computing, applications completely move from dedicated hardware/software at the enterprise or a specific datacenter into the 'cloud'.

In cloud computing, applications completely move from dedicated hardware/software at the enterprise or a specific datacenter into the 'cloud'.

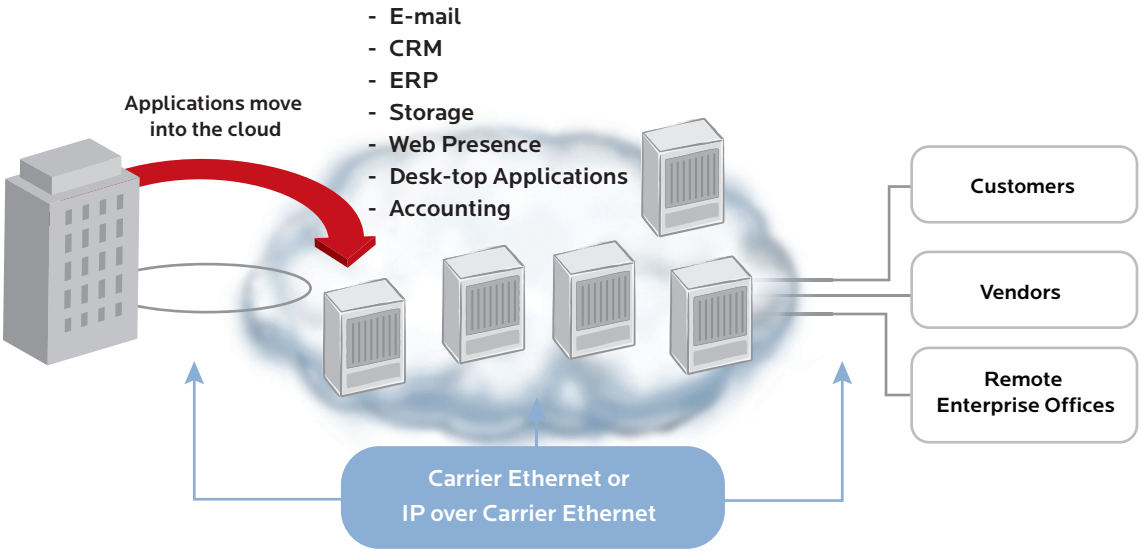

Figure 3. Cloud Architecture

On the surface, the cloud architecture (Figure 3) may not look that much different from applications in a traditional datacenter. The difference beneath the surface is there are no dedicated servers for each enterprise or even dedicated applications for each application. Applications can be instantiated on any server in the cloud, and can grow, migrate or replicate to other platforms as needed to enable scalability, optimization or resiliency.

Some companies like Google and Amazon are aggressively developing and delivering services for the complete end-to-end cloud market. These full service cloud infrastructure companies provide the physical servers and datacenters, the cloud operating system infrastructure, and applications. Others companies are targeting subsets of the cloud where they have a core competency.

The complete cloud solution is made up of the following components:

- Physical Server Infrastructure
- Operating System Foundation
- Application Layer
- Network

Virtualized applications on multiple servers that can be located anywhere in the cloud is a significant design change. Most applications were designed to run on a single server instance or at most a few well defined servers at once.

PHYSICAL SERVER INFRASTRUCTURE

The end game for cloud computing is for applications to be completely virtualized to run on any server in the cloud. Sometimes called Infrastructure-as-a-Service (IaaS) or utility computing, these can scale on demand, eliminating the need for dedicated servers per user or dedicated instances of software for each user.

This complete virtualization results in even more economy than the traditional datacenter model, where servers are dedicated to a customer application, and a separate instance of the software must be installed, managed, maintained and customized for each customer. Since average enterprise servers are only 5%-15% utilized, this enables a significant improvement in operating efficiency and cost.

The typical 100,000 sq/ft. datacenter costs $45M to $60M to construct. The power and cooling costs are enormous; datacenters are designed to support from 200-500 Watts of power per square foot which translates to 20-50 Megawatts of power per datacenter. Power is such a huge topic that access to cost-effective and scalable power is a major decision point on where datacenters are located. For example, Google located a cloud datacenter in The Dalles, Oregon to be next to a hydroelectric power plant on the Columbia River.

Datacenters are rated on their PUE (Power Use Effectiveness), which is a ratio of total power consumed to power actually use by the server equipment; the excess goes to cooling the equipment. A typical datacenter has a PUE of 2:1, meaning that for every 2 Megawatts of power consumed, 1 Megawatt is used by the computing equipment and 1 Megawatt is used for cooling and management. Cloud computing offers the promise of improving this ratio to 1.3:1, a 70% improvement in power efficiency. This is a significant improvement in the green factor - greener for the environment and more “green” for the bottom line.

OPERATING SYSTEM FOUNDATION

Virtualized applications on multiple servers that can be located anywhere in the cloud is a significant design change. Most applications were designed to run on a single server instance or at most a few well defined servers at once. Completely virtualized applications that can scale and replicate to new servers on the fly require a whole new operating system construct. Sometimes called Platform-as-a-Service (Paas), these solutions provide the foundation and application development environment to design, develop, test and deploy highly virtualized custom applications.

Notable examples include:

- Google App Engine (GAE)
- Amazon Web Services (AWS)
- Force.com

APPLICATION LAYER

Cloud enabled applications differ from the traditional enterprise IT model where the enterprise owns and manages a dedicated instance of the application. Software-as-a-Service (SaaS) enables complete virtualization and sharing of application processing and storage resources in a one-to-many environment. The primary advantages include:

- Pay per use or subscription based payment methods
- Rapid scalability
- Common version control - all clients are on same codebase, easy to manage software upgrades
- Economy of scale for software management and security
- No local client or server installation, access from any network connected device.

SaaS solutions include:

Customer Relationship Management (CRM) - Tools for managing customer interaction, sales, marketing, customer service and technical support. Example companies include Salesforce.com, SAP and Oracle.

Enterprise Resource Planning (ERP) - Tools for managing finance, accounting, human resources, manufacturing, sales and supply chain information across the entire organization. Example companies include Microsoft ERP, Oracle ERP, SAP and Ariba.

Security - As customer applications move into the cloud from the customer location, security functions which used to reside on the desktop or at the enterprise firewall move as well. Example companies which have cloud enabled security solutions include RSA, Symantec and IBM. Cloud solutions can provide traditional security with more flexibility and economy as well as new embedded security capabilities required when applications are distributed within the cloud.

Storage - Traditional Storage Area Networking (SAN) and Network Attached Storage (NAS) applications move into the cloud using general purpose servers and disk space. Storage becomes one more application that can reside anywhere in the cloud instead of a specialized application on specialized devices. This enables an order-of-magnitude reduction in the cost of storage while enabling more flexibility.

Desktop Applications - Tools that were traditionally dedicated on each desktop are being adapted to become solutions instantiated in the SaaS model. Examples include Google Apps, Webex and Quickbooks.

CLOUD NETWORKING

The connectivity between the cloud servers and cloud computing datacenters is a fundamental enabler of the cloud. Equally important is the access network that connects the enterprise and customer to the cloud.

Carrier Ethernet is ideally suited for both cloud access networking and cloud infrastructure networking.

This becomes evident when you consider the networking requirements of cloud computing:

Highly Scalable Bandwidth - Carrier Ethernet enables flows from 1Mb to 10GigE, with complete service differentiation and service assurance capabilities.

Service Differentiation - Robust Class-of-Service (CoS) and Quality-of-Service (QoS) are required to enable scaling of many customers and applications on a shared infrastructure. Carrier Ethernet solutions with flow level Class-of-Service and hierarchical Quality-of-Service enable guaranteed bandwidth per customer, differentiation between low and high priority applications, and differentiation between real-time nd non-real-time applications.

Performance Assurance - Customers and service providers need tools to guarantee end-to-end performance. As applications move into the cloud, customers require even move network performance visibility. Enterprises are entrusting their mission critical applications to someone outside of their firewall, and there are multiple players in the end-to-end service. Carrier Ethernet provides tools for end-end connectivity and performance monitoring with IEEE802.3ah, IEEE802.1ag, and ITU Y.1731.

Resiliency - Mission critical applications will not be trusted to best-effort connections. When enterprises move applications off-site into the cloud, they cannot tolerate loss of connectivity. Carrier Ethernet enables protection switching to a redundant path in the event of a fiber cut or node failure in less than 50ms using Link Aggregation or ITU G.8032 Ethernet Ring. Market leading solutions switch in 20ms.

Ethernet Economics - Ethernet is the most cost effective interface to the customer LAN and on the routers, switches and servers in the cloud computing datacenters. Carrier Ethernet empowers these devices to connect directly to the WAN using their native interfaces by providing all the performance, security and management features required in the WAN environment. Removing the interworking with legacy networking technologies simplifies the network, reduces CAPEX and OPEX.

Enterprises want solutions that are based on Carrier Ethernet due the superior cost, performance, reliability and scalability advantages. Network service providers have chosen Carrier Ethernet for the same reasons - it enables massive scaling of their network and services, provides integrated SLA (Service Level Agreement) management tools, high resiliency and a significant CAPEX and OPEX advantage over alternative technologies.

Connections within the cloud infrastructure include direct connections between data-centers and carrier-to-carrier interconnections. These are already 1GigE and 10GigE connections, and in many cases will grow to 40/100GigE in the future. Carrier Ethernet over direct fiber or as a wavelength on a DWDM solution is the optimal solution to meet the cost, performance and scalability for these connections. Carrier Ethernet service providers may use private network interconnection, or may use Carrier Ethernet Exchange points to interconnect with other carriers to provide the end-end core connectivity. A Carrier Ethernet Exchange provides a neutral interconnect point between service providers.

Cloud access networking connects enterprise offices to 'the cloud'. Access to the cloud can occur in a number of ways:

- **Direct connection** - Large enterprises may use a direct connection to a cloud-enabled datacenter. 1GigE or 10GigE Carrier Ethernet on a dedicated fiber or as a wavelength on a DWDM (Dense Wavelength Division Multiplexing) system provides resilient, scalable connectivity to the cloud datacenter. Ethernet over SONET/SDH may also be used.
- **Carrier Ethernet Exchange** - Carrier Ethernet Exchanges are not just for intra-cloud connections, they bring benefits to access as well. Enterprises benefit by having the larger geographic coverage and diversity options of all the exchange carrier member networks. In this way, an enterprise that does not have direct connectivity to a cloud location can realize 'direct' connectivity with access to the enterprise exchange provided by Carrier A which then interconnects via another carrier from the exchange node to cloud nodes.
- Ethernet Virtual Private Networking - Carrier Ethernet VPN services enable scalable, guaranteed bandwidth services over a switched Carrier network. This provides the same performance as a dedicated network, with the cost advantage of leveraging a shared network infrastructure from a Carrier Ethernet service provider. For example, an enterprise location that needs 20Mb of access to cloud services realizes the cost benefit of services provided over a multi-Gigabit or 10GigE network.
- **Public Internet Access** - Some enterprise locations may use Layer-3 VPN applications to enable secure access to cloud services. In these applications, the access to the Internet Service Provider (ISP) is optimized by using a Layer-2 Carrier Ethernet infrastructure.

Carrier Ethernet service providers may use private network inter-connection, or may use Carrier Ethernet Exchange points to interconnect with other carriers to provide the end-end core connectivity.

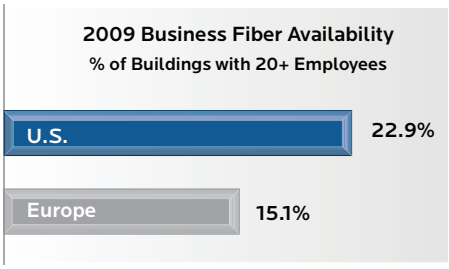

Source: Vertical Systems Group 2011

For all of these access methods, enterprises reap the cost and performance advantages of 1GigE and 10GigE native Ethernet infrastructure wherever fiber is available to the enterprise.

Carrier Ethernet is a smart way to leverage existing infrastructure as well. Even though the percentage of enterprises served by fiber has doubled in the last 5 years, there are still a large number of enterprise locations that do not have fiber connectivity. According to Vertical Systems group, only 15% of European and 23% of US businesses have fiber access. This leaves over 1 Million enterprise locations globally without direct fiber access. Ethernet over SONET/SDH or Ethernet over PDH (DS1/E1, DS3/E3) is a smart way to leverage existing infrastructure to provide ubiquitous service and stretch limited CAPEX dollars. Intelligent Carrier Ethernet Access Platforms provide the flow-level service differentiation and service assurance features to leverage existing transport. This enables consistent SLA capability and seamless end-end fault and performance management which are agnostic to the underlying transport network.

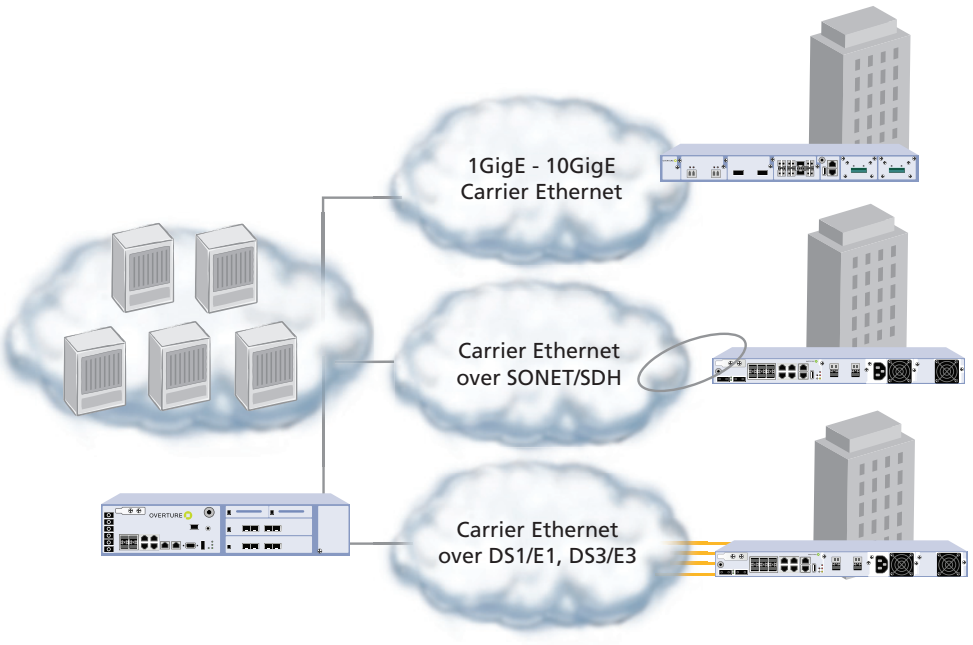

Figure 4. Carrier Ethernet access over any infrastructure

Many small-medium enterprise locations use legacy data transport solutions for their Public IP and Private IP applications. These include IP/PPP, IP/MLPPP and IP or Frame Relay over single DS1/E1 or bonded DS1/E1.

In many cases, these CPE solutions will stay in place for years until the customer is ready to transition to a new technology. This creates a challenge for the carrier who needs to convert his service edge switch/routers to scalable GigE interfaces and eliminate the expensive, processing intensive channelized interfaces.

A smart solution in these cases is to use an IP services interworking and aggregation solution that interworks between these legacy protocols and native Ethernet. This solution is inserted between the CPE and the cloud service edge, either at a remote POP or at the service hub site. This offloads the expensive, low-density channelized interfaces from the service edge switch/router, freeing it to scale using cost effective, high-density GigE interfaces. The CPE equipment still utilizes the existing interfaces without any CAPEX or operational changes required.

ACRONYMS

CCM	Continuity
ASP	Applications Service Provider
ATM	Asynchronous Transport Mode
FDDI	Fiber Distributed Data Interface
BTB	Business to Business
CoS	Class of Service
CRM	Customer Relationship Management
CTB	Customer to Business
ERP	Enterprise Resource Planning
GFP	Generic Frame Procedure
HVAC	Heating, Ventilation and Air Conditioning
IaaS	Infrastructure as a Service
ISP	Internet Service Provider
LAN	Local Area Network
LCAS	Link Capacity Adjustment Scheme
MLPPP	Multi-Link Point-to-Point Protocol
NAS	Network Attached Storage
PaaS	Platform as a Service
PUE	Power Use Effectiveness
QoS	Quality of Service
SaaS	Software as a Service
SAN	Storage Area Networking
SDH	Synchronous Digital Hierarchy
SLA	Service Level Agreement
SONET	Synchronous Optical Network
TDM	Time Domain Multiplexing
VCAT	Virtual Concatenation
VPN	Virtual Private Networking
WAN	Wide Area Network

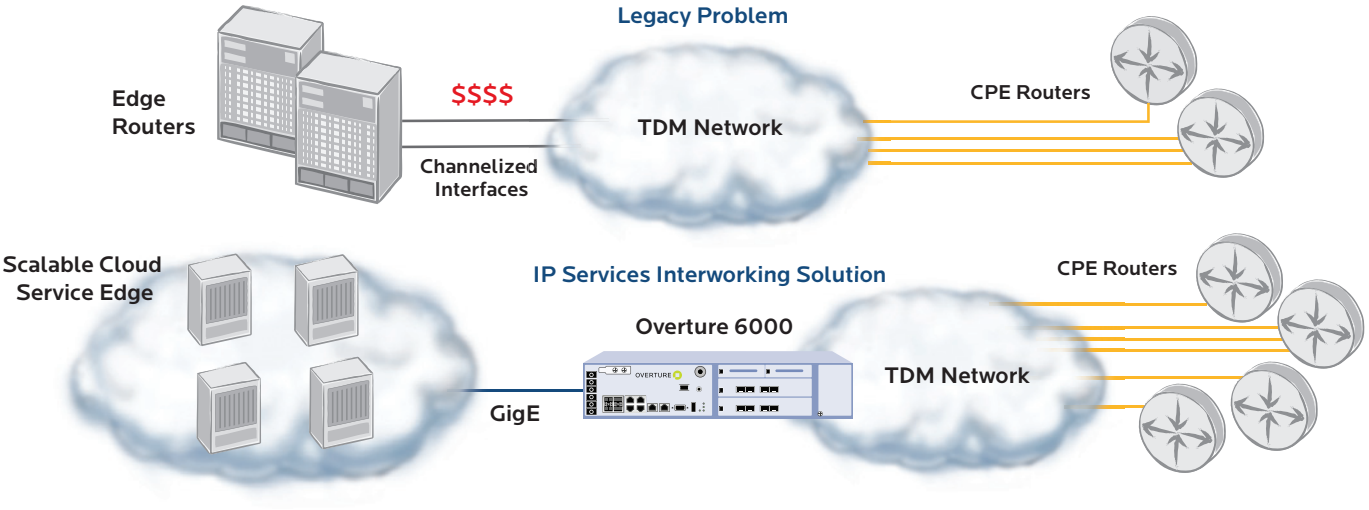

Figure 5. IP-services Interworking brings the Cloud to Legacy Locations

SUMMARY

Cloud networking is a fundamental component of cloud computing, there could be no cloud without a network to connect customers to the cloud, and to connect the servers and applications within the distributed cloud.

As applications move into the cloud, performance and reliability requirements increase. Enterprises are entrusting their mission critical applications to providers outside their fire-wall. The cloud is not one monolithic provider; it is made up of many networks and service providers which further increases the requirement for high performance networking with fine-grain service differentiation and assurance capabilities.

Cloud infrastructure providers leverage 1GigE and 10GigE Carrier Ethernet connections inside the cloud today, growing to 40/100GigE in the future. Enterprises and access service providers use 1GigE and 10GigE Carrier Ethernet where fiber is available to connect enterprise locations to the cloud. Ethernet over SONET/SDH or Ethernet over PDH are smart ways to leverage existing infrastructure to enable ubiquitous cloud services when native Ethernet over fiber or Ethernet radio are not available.

Carrier Ethernet is the optimal network solution to deliver the scalability, performance, reliability and economics required by cloud computing. Every network service provider that is delivering Carrier Ethernet services is either directly or indirectly a part of cloud computing services, and the growth trend is accelerating. For Carrier Ethernet service providers, a 'Cloudy' forecast is something to smile about.

ABOUT OVERTURE

Overture is the preferred Carrier Ethernet edge and aggregation partner to more than 450 service providers and enterprise customers worldwide. By providing the entrance to a better network, service providers can leverage Carrier Ethernet to multiply revenue and streamline operational costs by enabling high-capacity Ethernet services over any physical media, including fiber, copper and TDM. Overture's solutions are designed for reliability and ease of use, and arm customers to compete for demanding applications such as cloud computing and mobile communications that require greater bandwidth and smarter networks. Overture is headquartered in Research Triangle Park, NC, with a technology center in Richardson, Texas. For more information, visit **www.overturenetworks.com.**

OVERTURE

Overture Networks, Inc.
Research Triangle Park, NC
tel: 919.387.4100
www.overturenetworks.com